

Active Integrity Constraints for Multi-Context Systems^{*}

Luís Cruz-Filipe¹, Graça Gaspar², Isabel Nunes², and Peter Schneider-Kamp¹

¹ Dept. of Mathematics and Computer Science, University of Southern Denmark

² BioISI—Biosystems & Integrative Sciences Institute, Faculty of Sciences,
University of Lisbon

Abstract. We introduce a formalism to couple integrity constraints over general-purpose knowledge bases with actions that can be executed to restore consistency. This formalism generalizes active integrity constraints over databases. In the more general setting of multi-context systems, adding repair suggestions to integrity constraints allows defining simple iterative algorithms to find all possible grounded repairs – repairs for the global system that follow the suggestions given by the actions in the individual rules. We apply our methodology to ontologies, and show that it can express most relevant types of integrity constraints in this domain.

1 Introduction

Integrity constraints (ICs) for databases have been an important topic of research since the 1980s [1]. An early survey [26] already identified over 90 relevant types of integrity constraints. Since then, significant effort has been focused not only on identifying inconsistencies, but also on repairing inconsistent databases.

The same problem has been studied in other domains of knowledge representation. Integrity constraints for deductive databases [2] were also considered in the 1980s. More recently, interest for integrity constraints has arisen in the ontology domain, with several approaches on how to define them and how to check their satisfaction [14, 19, 21]. Given its challenges, the more complex problem of repairing inconsistent knowledge bases has not received as much attention.

In this paper, we address the problem of computing repairs by combining two ideas: clausal-form integrity constraints for multi-context systems (MCSs) [11] and active integrity constraints (AICs) for relational databases [16]. We demonstrate the expressiveness of our formalism and show how it can be used to compute repairs for inconsistent MCSs in general, and for ontologies, in particular.

Contribution. The main contribution of this paper is a notion of AIC for MCSs, which enables us to compute repairs for inconsistent MCSs automatically, requiring only decidability of entailment in the individual contexts. Particularized to

^{*} This work was supported by the Danish Council for Independent Research, Natural Sciences, grant DFF-1323-00247, and by FCT/MCTES/PIDDAC under centre grant to BioISI (Centre Reference: UID/MULTI/04046/2013).

ontologies, our framework is expressive enough to capture all types of integrity constraints identified as relevant in [14], as we exemplify in the text.

The step from ICs for MCSs to AICs for MCSs is inspired by the similar step in the database case [16]. However, we draw more significant benefits in this more general setting. AICs are ICs that also specify possible repair actions in their head. In the database case, every clausal IC can be transformed into an AIC automatically. The goal, though, is to restrict in order to establish preferences among different possible repairs. In the general case, such a transformation would require solving complex abduction problems [17].

Using AICs, we can automatically compute repairs for inconsistent MCSs, bypassing the need to solve such reasoning problems. The price to pay is the need to prove that an AIC is valid (Definition 4). The key observation here is that AICs should be written with a very clear semantic idea in mind, typically by an engineer with a deep knowledge of the underlying system, who should be able to show their validity formally. Thus, in practice, the complexity involved in computing each repair is moved to a one-time verification of validity of AICs.

Structure. We review previous work in Section 2, summarizing the key notions from [5, 8, 11]. Section 3 introduces AICs for MCSs, showing that they generalize the corresponding notion for relational databases, and studies their properties in general. Section 4 focuses on the case of ontologies and evaluates our formalism against the classes of integrity constraints identified in [14]. Section 5 discusses how algorithms to compute repairs in the database setting can be adapted to the general case of MCSs. We conclude in Section 6.

1.1 Related work

Database repairs. ICs for databases have been extensively studied throughout the last decades, and we restrict ourselves to works most directly related to ours.

Integrity constraints are typically grouped in different syntactic categories [26]. Many important classes can be expressed as first-order formulas, and can also be written in denial (clausal) form – the fragment expressible in our formalism.

Whenever an integrity constraint is violated, the database must be *repaired* to regain consistency. The problem of database repair is to determine whether such a transformation is possible, and many authors have invested in algorithms for computing database repairs efficiently. Typically, there are several possible ways of repairing an inconsistent database, and several criteria have been proposed to evaluate them. Minimality of change [13, 27] demands that the database be changed as little as possible, while the common-sense law of inertia [23] states that every change should have an underlying reason. While these criteria narrow down the possible database repairs, it is commonly accepted that human interaction is ultimately required to choose the “best” possible repair [25].

Active integrity constraints (AICs). The formalism of AICs, introduced in [16], addresses the issue of choosing among several possible repairs. An AIC specifies not only an integrity constraint, but it also gives indications on how inconsistent

databases can be repaired through the inclusion of *update actions*, which can be addition and removal of tuples from the database – a minimal set that can implement the three main operations of database updates [1].

The original, declarative, semantics of AICs defined *founded* repairs [5], in which every action is *supported*: it occurs in the head of a constraint that is violated if that action is not included. Despite this characterization, there are unnatural founded repairs where two actions mutually support each other, but do not have support from other actions. The same authors then proposed *justified* repairs [7], which however are not intuitive and pose further problems [9]. Furthermore, justified repairs are intrinsically linked to the syntactic structure of databases, and cannot be adapted to other knowledge representation formalisms.

Grounded repairs [8] form a middle ground between both semantics, requiring support for arbitrary subsets of the repair. They are grounded fixed points of the intuitive operation of “applying one action from the head of each AIC that is not satisfied”, which is in line with the intuitive motivation for studying AICs.

Founded and justified repairs can be computed via revision programming [7]. Alternatively, an operational semantics for AICs [9] was implemented for SQL databases [10]. There, repairs are leaves of particular trees, yielding a semantics equivalent to the declarative one when existence of a repair is an NP-complete problem. For grounded and justified repairs, where this existence problem is Σ_2^P -complete, the trees still contain all repairs, but may also include spurious leaves – requiring a post test that brings the overall complexity to the theoretical limit.

Multi-context systems (MCSs). MCSs, as defined in [3], can be informally described as collections of logic knowledge bases – the *contexts* – connected by Datalog-style *bridge rules*. Since their introduction, several variants of MCSs have been proposed that add to their potential fields of application. Relational MCSs [15] were proposed as a way to allow a formal first-order syntax, introducing variables and aggregate expressions in bridge rules, and extending the semantics of MCSs accordingly. Managed MCSs [4], which we describe in Section 2, further generalize MCSs by abstracting from the possible actions that change individual knowledge bases. Other variants, which are not directly relevant for this work, are discussed in [11]. A different line of research deals with repairing *logical* inconsistency of an MCS (non-existence of a model) [12].

ICs in ontologies. Integrating ICs with ontology-based systems poses several challenges, mainly due to the open-world assumption and the absence of the unique name assumption [14, 20, 22, 24]. In this context, ICs are conventionally modelled as T-Box axioms [19], but variants based on hybrid knowledge bases, auto-epistemic logic, modal logic, and grounded circumscription have recently been proposed. For an overview of these proposals see Section 2 in [21]. For details on how some of these can be expressed by ICs over MCSs, using a systematic interpretation of ontologies as MCSs, see Section 4.5 in [11]. The interpretation we use in Section 4 is a variant of the one presented therein.

2 Background

AICs for databases. Let Σ be a first-order signature without function symbols. A database is a set of ground atoms over Σ , and an update action is an expression of the form $+a$ or $-a$, where a is a ground atom over Σ . An *active integrity constraint* (AIC) over a database DB is a rule r of the form

$$p_1, \dots, p_m, \text{not } (p_{m+1}), \dots, \text{not } (p_\ell) \implies \alpha_1 \mid \dots \mid \alpha_k \quad (1)$$

where each p_i is an atom over the database's signature, every variable free in p_{m+1}, \dots, p_ℓ occurs in p_1, \dots, p_m , and each update action α_i is either $-p_j$ for some $1 \leq j \leq m$ or $+p_j$ for $m < j \leq \ell$.¹ The *body* of r is $\text{body}(r) = p_1, \dots, p_m, \text{not } (p_{m+1}), \dots, \text{not } (p_\ell)$, and the *head* of r is $\text{head}(r) = \alpha_1 \mid \dots \mid \alpha_k$.

If r is ground, then DB *satisfies* r , denoted $\text{DB} \models r$, if $\text{DB} \not\models p_i$ for some $1 \leq i \leq m$ or $\text{DB} \models p_i$ with $m < i \leq \ell$. In general, $\text{DB} \models r$ if DB satisfies all ground instances of r . Otherwise, r is *applicable* in DB [16]. If η is a set of AICs, then $\text{DB} \models \eta$ if $\text{DB} \models r$ for every $r \in \eta$.

A set of update actions \mathcal{U} is *consistent* if it does not contain both $+a$ and $-a$ for any ground atom a . Given a consistent \mathcal{U} , we write $\mathcal{U}(\text{DB})$ for the result of applying all actions in \mathcal{U} to DB, and say that \mathcal{U} is a *weak repair* for $\langle \text{DB}, \eta \rangle$ if: (i) every action in \mathcal{U} changes DB and (ii) $\mathcal{U}(\text{DB}) \models \eta$. \mathcal{U} is a *repair* if $\mathcal{V}(\text{DB}) \not\models \eta$ for every $\mathcal{V} \subsetneq \mathcal{U}$ [5], and \mathcal{U} is *grounded* if, for every $\mathcal{V} \subsetneq \mathcal{U}$, there exists a ground instance r of a rule in η such that $\mathcal{V}(\text{DB}) \not\models r$ and $\text{head}(r) \cap (\mathcal{U} \setminus \mathcal{V}) \neq \emptyset$ [8].

Multi-Context Systems. We now describe the variant of multi-context systems we use: *managed multi-context systems* (also abbreviated to MCSs) [4].

A *relational logic* L is a tuple $\langle \text{KB}, \text{BS}, \text{ACC}, \Sigma \rangle$, where KB is the set of well-formed knowledge bases of L (sets of well-formed formulas), BS is a set of possible belief sets (candidate models), $\text{ACC} : \text{KB} \rightarrow 2^{\text{BS}}$ is a function assigning to each knowledge base a set of acceptable belief sets (its models), and Σ is a signature generating first-order sublanguages of $\bigcup \text{KB}$ and $\bigcup \text{BS}$.

A *managed multi-context system* is a collection of managed contexts $\{C_i\}_{i=1}^n$, with each $C_i = \langle L_i, \text{kb}_i, \text{br}_i, D_i, \text{OP}_i, \text{mng}_i \rangle$ where: $L_i = \langle \text{KB}_i, \text{BS}_i, \text{ACC}_i, \Sigma_i \rangle$ is a relational logic; $\text{kb}_i \in \text{KB}_i$; D_i (the *import domain*) is a set of constants from Σ_i ; OP_i is a set of operation names; $\text{mng}_i : \wp(\text{OP}_i \times \bigcup \text{KB}_i) \times \text{KB}_i \rightarrow \text{KB}_i$ is a *management function*; and br_i is a set of *managed bridge rules*, with the form

$$(i : o(p)) \leftarrow (i_1 : p_1), \dots, (i_q : p_q), \text{not } (i_{q+1} : p_{q+1}), \dots, \text{not } (i_m : p_m) \quad (2)$$

such that $o \in \text{OP}_i$, $p \in \bigcup \text{KB}_i$, $1 \leq i, i_j \leq n$, and each p_j is a belief² of L_{C_j} .

Intuitively, kb_i is the knowledge base of context C_i and OP_i are the names of the operations that can be applied to change it. The management function

¹ In [16], existentially quantified variables can also occur in negative literals. This was not discussed in subsequent work, and we ignore it for simplicity of presentation.

² Technically, P_p is a *relational element* of C_{i_p} : it can include variables, which when instantiated yield elements of $\bigcup \text{BS}_{i_p}$ – see [4] for details.

defines the semantics of these operations: $\text{mng}_i(O, \text{kb})$ is the result of applying the operations in O to kb . Bridge rules govern the interaction between contexts.³

A *belief state* for an MCS $M = \{C_i\}_{i=1}^n$ is a set $S = \{S_i\}_{i=1}^n$ such that each $S_i \in \text{BS}_i$. A ground instance of bridge rule (2) is *applicable* in S if $p_i \in S_i$ for $1 \leq i \leq q$ and $p_i \notin S_i$ for $q < i \leq m$; the variables in the rule can only be instantiated by elements of the import domain D_i . A belief state is an *equilibrium* for M if it is stable under application of all bridge rules, i.e.:

$$S_i \in \text{ACC}_i(\text{mng}_i(\{\text{head}(r) \mid r \in \text{br}_i \text{ applicable in } S\}, \text{kb}_i))$$

In general, M can have zero, one or several equilibria; if at least one exists, then M is *logically consistent*. We present examples of MCSs in the next sections.

Integrity constraints for general-purpose knowledge bases. ICs for MCSs [11] generalize clausal ICs to a generic framework for reasoning systems – covering not only relational databases, but also deductive databases, peer-to-peer systems and ontologies, among others. Syntactically, ICs are bridge rules with empty head, forming an added layer on top of an MCS that does not affect its semantics.

As MCSs may have several equilibria, satisfaction of a set of ICs η can be *weak* – there is an equilibrium satisfying all rules in η – or *strong* – all equilibria satisfy all rules in η . In order to avoid vacuous quantifications, strong satisfaction only holds for logically consistent MCSs. In general these properties are undecidable [11], but if entailment in every context is decidable then satisfaction of a set of ICs is in most cases as hard as the hardest entailment decision problem.

In this paper, we do not explicitly mention the set of ICs when clear from the context. Moreover, our development applies both to weak and strong satisfaction, and we simply say that an MCS is *consistent* if it satisfies the given set of ICs. We explicitly write “logical consistency” for existence of an equilibrium.

3 Active Integrity Constraints

We begin by defining active integrity constraints over multi-context systems.

Definition 1. *An AIC over an MCS $M = \{C_i\}_{i=1}^n$ is a rule r of the form*

$$(i_1 : P_1), \dots, (i_m : P_m), \text{not } (i_{m+1} : P_{m+1}), \dots, \text{not } (i_\ell : P_\ell) \\ \implies (j_1 : \alpha_1) \mid \dots \mid (j_k : \alpha_k) \quad (3)$$

where $1 \leq i_p, j_q \leq n$, each P_p is a belief in C_{i_p} , each update action $\alpha_q \in \text{OP}_{j_q} \times \bigcup \text{KB}_{j_q}$, and all variables in P_{m+1}, \dots, P_ℓ occur in P_1, \dots, P_m .

This definition follows the one for databases (1), and we define *body* and *head* of r similarly. Equation (3) also generalizes ICs for MCSs: each AIC corresponds

³ For the sake of presentation, we simplified the management function, which in the original work is allowed to return several possible effects for each action.

to an IC by ignoring its head, immediately yielding notions of weak and strong satisfaction for an AIC. We also say that r is *applicable* to an MCS M if $M \not\models r$. Intuitively, in this case M should be repaired by applying actions in $\text{head}(r)$.

The reasoning capabilities of MCSs dictate that we cannot restrict the actions in the head of an AIC syntactically (as in the database world, see Section 2). We thus relax this requirement by only demanding that the actions are capable of solving the inconsistency. It is also not reasonable to require that every action in $\text{head}(r)$ be able to solve every inconsistency detected by $\text{body}(r)$: since inconsistencies may be triggered by derived information, they may have different origins, and the different actions may be solutions for those different causes.

We are interested in sets of update actions that are applied simultaneously, i.e. the order in which actions are executed should be irrelevant. This corresponds to the consistency requirement usually considered in databases.

Definition 2. Let $M = \{C_i\}_{i=1}^n$ be an MCS, \mathcal{U} be a finite set of update actions, and \mathcal{U}_i be the set of actions in \mathcal{U} affecting C_i .

\mathcal{U}_i is consistent w.r.t. kb_i if, for every permutation $\alpha_1, \dots, \alpha_k$ of the elements of \mathcal{U}_i , $\text{mng}_i(\mathcal{U}_i, \text{kb}_i) = \text{mng}_i(\alpha_1, \text{mng}_i(\dots, \text{mng}_i(\alpha_k, \text{kb}_i) \dots))$. \mathcal{U} is consistent w.r.t. M if each \mathcal{U}_i is consistent w.r.t. kb_i , and in this case we write $\mathcal{U}(M)$ for the result of applying each \mathcal{U}_i to each kb_i .

Example 1. We consider a concrete toy example of a deductive database with two unary base relations p and q , a view consisting of a relation r such that $r(x) \leftrightarrow p(x) \vee q(x)$, and the integrity constraint $\neg r(a)$.

We formalize this as an MCS $M = \langle C_E, C_I \rangle$ where C_E is an extensional database including predicates p and q (but not r), C_I is the view context including predicate r (but not p or q), and they are connected by the bridge rules

$$(I : r(X)) \leftarrow (E : p(X)) \quad (I : r(X)) \leftarrow (E : q(X)).$$

Furthermore, mng_E allows addition and removal of any tuples to C_E , using operations **add** and **del**, while mng_I does not allow any changes. (See [11] for details of this construction.)

From the structure of M , we know that $r(a)$ can only arise as a deduction from $p(a)$ or $q(a)$ (or both), so it makes sense to write an AIC

$$(I : r(a)) \implies (E : \text{del}(p(a))) \mid (E : \text{del}(q(a))).$$

The actions on the head of this AIC solve the problem in all future states of M , since C_I cannot change. However, restoring consistency may require performing both actions (if the database contains both $p(a)$ and $q(a)$).

This example also illustrates an important point: repair actions are written with a particular structure of the MCS in mind.

Definition 3. The set of variants to an MCS M , denoted $\text{vrt}(M)$, is

$$\text{vrt}(M) = \{\mathcal{U}(M) \mid \mathcal{U} \text{ is a finite set of update actions over } M\}.$$

Restrictions on the actions in the head of AICs only range over $\text{vrt}(M)$, which contains all possible future evolutions of M .

Definition 4. An AIC r of the form (3) is valid w.r.t. an MCS M if:

- for every logically consistent $M' \in \text{vrt}(M)$ such that $M' \not\models r$, there is $\mathcal{U} \subseteq \text{head}(r)$ with $\mathcal{U}(M') \models r$;
- for every $\alpha \in \text{head}(r)$, there is $M' \in \text{vrt}(M)$ with $M' \not\models r$ and $\alpha(M') \models r$.

These conditions require that the set of suggested actions be complete (it can solve all inconsistencies) and that it does not contain useless actions.

Example 2. The AIC in Example 1 is valid: the only possible changes to M are in kb_E , which only contains information about p and q , thus, in any element of $\text{vrt}(M)$ the only way to derive $r(a)$ is still from either $p(a)$ or $q(a)$. The second condition follows by considering M' with $\text{kb}_E = \{p(a)\}$ and $\text{kb}_E = \{q(a)\}$.

Proposition 1. Deciding whether an AIC is valid is in general undecidable.

Proof (sketch). Let L be a logic with an undecidable entailment problem, C be a context over L with $\text{add} \in \text{OP}_C$ such that $\text{mng}_C(\text{add}(\varphi), \Gamma) = \Gamma \cup \{\varphi\}$, and $M = \{C\}$. Assume also that $\text{vrt}(M)$ includes all knowledge bases over L . Then $(C : \neg B) \implies (C : \text{add}(A))$ is valid iff $A \models_L B$. \square

In practice, proving validity of AICs should not pose a problem: AICs are written by humans with a very precise semantic motivation in mind, and this means that the conditions in Definition 4 should be simple for a human to prove.

We now show that the framework we propose generalizes the database case. A database DB can be seen as an MCS $M(\text{DB})$, defined as having a single context over first-order logic, whose knowledge base is DB , with management function allowing addition (+) or removal (–) of facts, and where the only set of beliefs admissible w.r.t. a given database is the set of literals that are true in that database (see [11] for a detailed definition).

Proposition 2. Every AIC over a database DB yields a valid AIC over $M(\text{DB})$.

Proof (sketch). We write a generic AIC over a database (1) as the AIC

$$(1 : p_1), \dots, (1 : p_m), \text{not } (1 : p_{m+1}), \dots, \text{not } (1 : p_\ell) \implies (1 : \alpha_1) \mid \dots \mid (1 : \alpha_k)$$

over $M(\text{DB})$. If DB does not satisfy the body of (1), then it can always be repaired by performing exactly one of the actions in its head [6], establishing both conditions for validity. \square

Definition 5. Let $M = \{C_i\}_{i=1}^n$ be an MCS, η be a set of AICs over M and \mathcal{U} be a finite set of update actions. \mathcal{U} is a weak repair for $\langle M, \eta \rangle$ if \mathcal{U} is consistent w.r.t. M and $\mathcal{U}(M) \models \eta$. Furthermore, \mathcal{U} is grounded if: for every $\mathcal{V} \subsetneq \mathcal{U}$, there is an AIC $r \in \eta$ such that $\mathcal{V}(M) \not\models r$ and $\text{head}(r) \cap (\mathcal{U} \setminus \mathcal{V}) \neq \emptyset$.

The definitions of weak and grounded repair directly correspond to those for the database case (Section 2). The notion of grounded repair implies, in particular, minimality under inclusion [8].

4 Application: the Case of Ontologies

This section is devoted to examples illustrating how our framework can be applied to the particular case of integrity constraints over ontologies.

Previous work [3, 11] shows how to view an ontology as a context of an MCS. In the present work, we refine this interpretation by representing an ontology as *two* contexts: one for the A-Box, one for the T-Box, connected by bridge rules that port every instance from the former into the latter. (This is reminiscent of how deductive databases are encoded in MCSs, see [11].) This finer encoding allows us, in particular, to reason about asserted instances (which are given in the A-Box) and those that are derived using the axioms (see Example 5).

We further assume that the A-Box only contains instances of atomic concepts or roles ($C(t)$ or $R(t, t')$). This option does not restrict the expressive power of the ontology, but it helps structure AICs: to include instance axioms about e.g. $C \sqcup D$, one instead defines a new concept $E = C \sqcup D$ in the T-Box and includes instance axioms about E in the A-Box (see also Example 7).

Definition 6. *A description logic \mathcal{L} is represented as the relational logic $L_{\mathcal{L}} = \langle \text{KB}_{\mathcal{L}}, \text{BS}_{\mathcal{L}}, \text{ACC}_{\mathcal{L}}, \Sigma_L \rangle$, where:*

- $\text{KB}_{\mathcal{L}}$ contains all well-formed knowledge bases of \mathcal{L} ;
- $\text{BS}_{\mathcal{L}}$ contains all sets of queries in the language of \mathcal{L} ;
- $\text{ACC}_{\mathcal{L}}(\text{kb})$ is the singleton set containing the set of queries to which kb answers “Yes”.
- Σ_L is the first-order signature underlying \mathcal{L} .

An ontology $\mathcal{O} = \langle T, A \rangle$ based on \mathcal{L} induces the multi-context system $M(\mathcal{O}) = \langle \text{Ctx}(T), \text{Ctx}(A) \rangle$ where $\text{Ctx}(T) = \langle L_{\mathcal{L}}, T, \text{br}_T, \Sigma_0, \emptyset, \emptyset \rangle$ with

- br_T contains all rules of the form $(T : C)(X) \leftarrow (A : C)(X)$ where C is a concept, and $(T : R)(X, Y) \leftarrow (A : R)(X, Y)$ where R is a role;
- Σ_0 is the set of constants in $\Sigma_{\mathcal{L}}$;

and $\text{Ctx}(A) = \langle L_{\mathcal{L}}, A, \emptyset, \Sigma_0, \text{OP}, \text{mng} \rangle$ where OP and mng are the set of allowed update operation names and their definition.

The management function does not allow changes to the T-Box; the particular operations in the A-Box depend on the concrete ontology. This is in line with our motivation that writing AICs requires knowledge of the system’s deductive abilities (expressed by the T-Box), which should not change.

We now evaluate the expressivity of our development by showing how to formalize several types of ICs over ontologies. We follow the classification in Section 4.5 of [14], which describes families of ICs determined by OWL engineers and ontologists as the most interesting, as well as other types of ICs considered in the scientific literature. Several classes of ICs are syntactically similar, so we do not include examples for all categories in [14], but explain in the text how the missing ones can be treated.

Most of our examples are adapted from [14], which frames them in a variant of the Lehigh University Benchmark [18], an ontology designed with the goal of providing a realistic scenario for testing. This ontology considers concepts `student`, `gradStudent`, `class` and `email`, and roles `hasEmail`, `enrolled` and `webEnrolled`. Our semantics is: `class` is a concept including all classes of a common course; `enrolled(c, s)` holds if student `s` is enrolled in course `s`; and `webEnrolled` holds if the student is furthermore to be contacted only electronically.⁴ The actual contents of the A-Box are immaterial for our presentation, and we restrict ourselves to the fragment of the T-Box containing the following axioms.

$$\begin{array}{ll}
 \text{gradStudent} \sqsubseteq \text{student} & \exists \text{enrolled.student} \sqsubseteq \text{class} \\
 \text{webEnrolled} \sqsubseteq \text{enrolled} & \exists \text{hasEmail.email} \sqsubseteq \text{student} \\
 & \exists \text{webEnrolled}^R.\text{class} \sqsubseteq \exists \text{hasEmail}
 \end{array}$$

4.1 Functional dependencies

Functional dependencies are one of the most frequently occurring families of ICs: requirements that certain relations be functional on one argument. In our example, this applies to `hasEmail`: two distinct students cannot have the same e-mail.

Since ontologies do not have the Unique Name Assumption, we cannot distinguish individuals by checking name equality (as in databases), but must query the ontology instead. Furthermore, while in the database world such violations can only be repaired by removing one of the offending instances, in ontologies, we can also add the information that two individuals are the same.

Example 3. Suppose that the management function includes operations `add` and `del` to add or remove a particular instance from the A-Box, as well as `assertEqual`, establishing equality of two individuals. Under these assumptions, we can express functionality of e-mail as the following AIC.

$$\begin{aligned}
 & (A : \text{hasEmail}(X, Z)), (A : \text{hasEmail}(Y, Z)), \text{not } (T : (X = Y)) \\
 & \implies (A : \text{del}(\text{hasEmail}(X, Z))) \mid (A : \text{assert}(X = Y)) \quad (4)
 \end{aligned}$$

Observe that, if T explicitly proves that $X \neq Y$, then only the first action can be used, as asserting equality between X and Y would lead to an inconsistency. However, if this is not the case then the second action is also a repair possibility, and hence this AIC is valid. There are several possibilities for the implementation of `assert`: it can add the equality $X = Y$ to the A-Box, but it can also syntactically replace every occurrence of one of them for the other.

Several other types of dependencies (e.g. key constraints, uniqueness constraints, functionality constraints) are expressed by similar formulas. Likewise, max-cardinality constraints can be represented as AICs with similar types of actions in the head (deleting some instances or unifying some individuals).

⁴ This semantics is slightly changed from that of [14], in order to make some aspects of our example more realistic.

4.2 Property domain constraints

This family of ICs specifies that the domain of a role should be a subset of a particular concept. In case such a constraint is violated, the offending element has to be added as an instance of that concept. The treatment of these ICs is thus very similar to the database case.

Example 4. To model that only students can be enrolled in courses, we write the following AIC.

$$(T : \text{enrolled}(X, Y), \text{not } (T : \text{student}(Y))) \implies (A : \text{add}(\text{student}(Y))) \quad (5)$$

We could also add the action $(A : \text{del}(\text{enrolled}(X, Y)))$ to the head of this AIC; note that it would only restore consistency in the case where this fact is explicitly stated in the A-Box and not otherwise derivable. Property range constraints (restricting the range of a role) can be similarly treated.

4.3 Specific type constraints

In many applications, it is interesting to minimize redundancy in the A-Box. In particular, in the presence of inclusion axioms, it is often desirable only to include instances pertaining to the most specific type class of each individual.

Example 5. Since $\text{gradStudent} \sqsubseteq \text{student}$, we guarantee that the A-Box only contains instances of the most specific class a student belongs to by writing:

$$(A : \text{gradStudent}(X), (A : \text{student}(X))) \implies (A : \text{del}(\text{student}(X))) \quad (6)$$

Thus, if the A-Box contains e.g. $\text{student}(\text{john})$ and $\text{gradStudent}(\text{john})$, then the axiom $\text{student}(\text{john})$ will be removed. The system will still be able to derive $\text{student}(\text{john})$, but only in context C_T (using the information in the T-Box). The separation of the A-Box and T-Box in different contexts is essential to express this integrity constraint in our formalism. Constraints that distinguish between assertions explicitly stated in the A-Box and derived ones have been considered e.g. in [22].

4.4 Min-cardinality constraints

We now consider a more interesting type of ICs: min-cardinality constraints. Inconsistencies arising from the violation of such constraints are hard to repair automatically, as such a repair requires “guessing” which instances to add. Using AICs and adequate management functions, we can even specify the construction of “default” values that may depend on the actual ontology.

Example 6. We want to express that each class must have a minimum of 10 students. Classes with less enrolled students should be closed, and those students moved to the smallest remaining class using an operation `redistribute`.

$$(T : (\leq 10.\text{enrolled})(X)) \implies (A : \text{redistribute}(\neg\text{class}(X))) \quad (7)$$

For this AIC to be valid, `redistribute` must check whether students are `enrolled` or `webEnrolled` and change the appropriate instance in the A-Box. This also uses the knowledge that instances of `enrolled` cannot be derived in other ways.

A similar kind of constraints are totality constraints, which require that a role be total on one of its arguments. In our example, we could require every student to be enrolled in some class, and use an adequate management function to add non-enrolled students to e.g. the smallest class.

4.5 Missing property value constraints

We now turn our attention to a kind of ICs that is also very common in ontologies: disallowing unnamed individuals for particular properties [22].

Example 7. Our ontology specifies that all students that are `webEnrolled` in a class must have an e-mail address. However, for the purpose of contacting these individuals, this e-mail address must be explicitly provided. We address this issue with the following AIC.

$$(T : (\exists \text{hasEmail})(X)), \text{ not } (T : \text{hasEmail}(X, Y)) \\ \implies (A : \text{unregister}(\neg \exists \text{webEnrolled}^R(X))) \quad (8)$$

Here, `unregister` replaces the axiom `webEnrolled(X)` with `enrolled(X)`, as it makes sense to keep the student enrolled in the course. Validity of this AIC follows from observing that the only possible ways to derive $\exists \text{hasEmail}(X)$ are either from an explicit assertion `hasEmail(X, Y)` or indirectly from `webEnrolled(Z, X)`.

This example also justifies our requirement that the A-Box can only contain instances of atomic concepts or roles. If the A-Box were allowed to contain e.g. $\exists \text{hasEmail}(\text{john})$, then AIC (8) would no longer be valid. By restricting to atomic concepts, the only way to perform a similar change would be by defining a new concept as equivalent to $\exists \text{hasEmail}$ – and this information would be present in the T-Box, making it clear that AICs should consider it.

4.6 Managing unnamed individuals

Finally, we illustrate how we can write AICs in different ways to control whether they range over all individuals of a certain class, or only over named ones.

Example 8. For ecological reasons, we want all students with an e-mail address to be enrolled in the web version of courses. We can write this as follows.

$$(T : (\text{hasEmail})(Y, Z)), (T : \text{enrolled}(X, Y)), \text{ not } (T : \text{webEnrolled}(X, Y)) \\ \implies (A : \text{webEnroll}(\text{webEnrolled}(X, Y))) \quad (9)$$

Operation `webEnroll` will replace `enrolled(X, Y)` with `webEnrolled(X, Y)`, dually to `unregister` in the previous example.

Alternatively, we could consider writing

$$(T : (\exists \text{hasEmail})(Y)), (T : \text{enrolled}(X, Y)), \text{not } (T : \text{webEnrolled}(X, Y)) \\ \implies (A : \text{webEnroll}(\text{webEnrolled}(X, Y))) \quad (10)$$

In this particular context, this formulation is undesirable, as it will also affect individuals who do not have a known e-mail address. By writing an explicit variable in the first query of the body, as in (9), we guarantee that we only affect those individuals whose e-mail address is known.

Similar considerations about the two possible ways to formulate this type of ICs can be found in [22].

5 Computing Repairs

In [9], we showed how to use active integrity constraints to compute repairs for inconsistent databases, by using the actions in the head of unsatisfied AICs to build a *repair tree* whose leaves were the repairs. We showed how the construction of the tree could be adapted to the different types of repairs considered originally in [7]; in particular, for the case of grounded repairs (which is the one we are interested in this work), it is enough to expand each node with the actions in the heads of the AICs that are not satisfied in that node.

We adapt this construction to the framework of AICs over MCSs. As we will see, the algorithms have to be adapted to this more general scenario, but we can still construct all grounded repairs for a given (inconsistent) MCS automatically, as long as entailment in all contexts is decidable.

Definition 7. *Let M be an MCS and η be a set of integrity constraints over M . The repair tree for $\langle M, \eta \rangle$, $\mathcal{T}_{\langle M, \eta \rangle}$, is defined as follows.*

- Each node is a set of update actions.
- A node n is consistent if: (i) $n(M)$ is logically consistent and (ii) if n' is the parent of n , then n is a consistent set of update actions w.r.t. $n'(M)$.
- Each edge is labeled with a closed instance of a rule.
- The root of the tree is the empty set \emptyset .
- For each consistent node n and rule r , if $n(M) \not\models r$ then $n' = n \cup \mathcal{U}$ is a child of n if (i) $\mathcal{U} \subseteq \text{head}(r)$, (ii) $n'(M) \models r$ and (iii) if $\mathcal{U}' \subseteq \mathcal{U}$ then $(n \cup \mathcal{U}')(M) \not\models r$.

In the database case [9], it is straightforward to show that repair trees are finite, since the syntactic restrictions on database AICs guarantee that each rule can only be applied at most once in every branch. In the general MCS case, this is not true, as the following example shows.

Example 9. Consider an ontology (represented as an MCS as in Section 3) with four concepts B_1 , B_2 , B_3 and D . The T-Box contains axioms

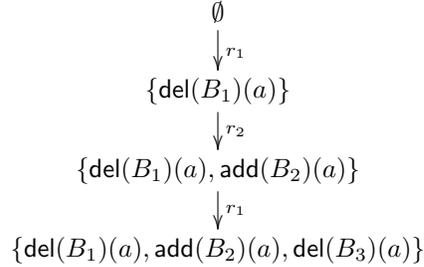
$$B_1 \sqsubseteq D \quad \text{and} \quad B_2 \sqcap B_3 \sqsubseteq D$$

and the A-Box is $\{B_1(a), B_3(a)\}$. Furthermore, we have integrity constraints

$$(T : D)(a) \implies (A : \text{del}(B_1)(a)) \mid (A : \text{del}(B_3)(a)) \quad (r_1)$$

$$\text{not } (T : B_1)(a), \text{not } (T : B_2)(a) \implies (A : \text{add}(B_2)(a)) \quad (r_2)$$

Following this construction, we obtain the tree on the right, and its leaf is a grounded repair.



Lemma 1. $\mathcal{T}_{\langle M, \eta \rangle}$ is finite.

Proof. By definition, every node of $\mathcal{T}_{\langle M, \eta \rangle}$ has a finite number of descendants, since there are only finitely many ground instances of AICs with a finite number of actions in each one's head. By construction, in every branch the labels of the nodes form an increasing sequence (w.r.t. set inclusion), and each node is again a subset of the (finite) set of all actions in the heads of all rules. Therefore, $\mathcal{T}_{\langle M, \eta \rangle}$ has finite depth and finite degree, hence it is finite. \square

Lemma 2. Every grounded repair for $\langle M, \eta \rangle$ is a leaf of $\mathcal{T}_{\langle M, \eta \rangle}$.

Proof. Let \mathcal{U} be a grounded repair for M and η . By definition of grounded repair, if $\mathcal{U}' \subseteq \mathcal{U}$ then there is a ground instance r of an AIC such that: there exists $\mathcal{V} \subseteq \text{head}(r) \cap \mathcal{U}$ such that $(\mathcal{U}' \cup \mathcal{V})(M) \models r$. This directly yields a branch of the repair tree ending at \mathcal{U} . \square

(This is essentially the same argument for showing that, in the database case, grounded repairs are well-founded, see [8].)

$\mathcal{T}_{\langle M, \eta \rangle}$ is constructed as the well-founded repair tree in the database case [9].⁵ In both cases, this tree may, in general, contain leaves that are not grounded repairs [8]. Under the assumption that $P \neq NP$, this cannot be avoided, since existence of grounded repairs for databases is a Σ_2^P -complete problem [8].

Complexity. The proof of Lemma 1 shows that the depth of $\mathcal{T}_{\langle M, \eta \rangle}$ is polynomial in the size of the grounded instances of η . Therefore, given an oracle that decides whether an MCS satisfies a set of AICs, the problem of existence of a grounded repair for $\langle M, \eta \rangle$ is Σ_2^P -complete: $\mathcal{T}_{\langle M, \eta \rangle}$ can be built in non-deterministic polynomial time (guessing which rule to apply at each node and using the oracle to decide whether the descendant is a leaf), and the validation step can be done in co-NP time (if the leaf is not a grounded repair, then we guess the subset that violates the definition and use the oracle to confirm this).

⁵ There is also a notion of repair tree for databases in [9], but it relies on the ability of inferring heads of AICs automatically, which does not exist in the MCS setting.

6 Discussion and Conclusions

Validity. At the end of Example 1, we pointed out that restoring consistent w.r.t. an AIC r may require applying several actions in $\text{head}(r)$. This suggests allowing sets of actions (rather than actions) in the heads of AICs. Besides increasing the complexity of our development, it is not clear that this change would bring significant benefits. In terms of computing repairs, we already cover those cases, since we add sets of actions when going from a node to its descendents. Also, it is not clear that there exists a situation when *every* possible inconsistent MCS requires a set of actions to repair.

One could also remove the second condition of validity of an AIC, i.e. allow the actions in the head to be insufficient to restore consistency of some MCSs. This would remove some burden from the programmer who has to specify the AICs, and would not affect the performance of the algorithms in Section 5. However, it would contradict the original motivation for AICs [16]: that the actions in the head of a rule should provide the means for restoring consistency.

Variants of AICs. The authors of [16] also considered *conditioned active integrity constraints*, where the actions on the head of AICs are guarded by additional conditions that have to be satisfied. In their setting, conditioned AICs do not add expressive power to the formalism, as they can be split into several unconditioned AICs (with more specific bodies) preserving the notions of consistency and repairs. In our setting, this transformation is not possible, and it would thus be interesting to study conditioned active integrity constraints over multi-context systems. However, we point out that the management function *can* use information about the actual knowledge bases in its implementation, so some conditions can actually be expressed in our setting (see Example 6).

Conclusion. We proposed active integrity constraints for multi-context systems and showed that, using them, we can compute grounded repairs for inconsistent MCSs automatically. Although validity of AICs is in general undecidable, we showed that we can cover the most common types of ICs in our framework.

References

1. S. Abiteboul. Updates, a new frontier. In *ICDT*, volume 326 of *LNCS*, pages 1–18. Springer, 1988.
2. P. Asirelli, M. de Santis, and M. Martelli. Integrity constraints for logic databases. *J. Log. Program.*, 2(3):221–232, 1985.
3. G. Brewka and T. Eiter. Equilibria in heterogeneous nonmonotonic multi-context systems. In *AAAI*, pages 385–390. AAAI Press, 2007.
4. G. Brewka, T. Eiter, M. Fink, and A. Weinzierl. Managed multi-context systems. In *IJCAI*, pages 786–791. IJCAI/AAAI, 2011.
5. L. Caroprese, S. Greco, C. Sirangelo, and E. Zumpano. Declarative semantics of production rules for integrity maintenance. In *ICLP*, volume 4079 of *LNCS*, pages 26–40. Springer, 2006.

6. L. Caroprese and M. Truszczyński. Declarative semantics for active integrity constraints. In *ICLP*, volume 5366 of *LNCS*, pages 269–283. Springer, 2008.
7. L. Caroprese and M. Truszczyński. Active integrity constraints and revision programming. *Theory Pract. Log. Program.*, 11(6):905–952, November 2011.
8. L. Cruz-Filipe. Grounded fixpoints and active integrity constraints. In *ICLP*, OASICS. Dagstuhl, 2016. Accepted for publication.
9. L. Cruz-Filipe, P. Engrácia, G. Gaspar, and I. Nunes. Computing repairs from active integrity constraints. In *TASE*, pages 183–190. IEEE, 2013.
10. L. Cruz-Filipe, M. Franz, A. Hakhverdyan, M. Ludovico, I. Nunes, and P. Schneider-Kamp. repAIrC: A tool for ensuring data consistency by means of active integrity constraints. In *KMIS*, pages 17–26. SciTePress, 2015.
11. L. Cruz-Filipe, I. Nunes, and P. Schneider-Kamp. Integrity constraints for general-purpose knowledge bases. In *FoIKS*, volume 9616 of *LNCS*, pages 235–254. Springer, 2016.
12. T. Eiter, M. Fink, G. Ianni, and P. Schüller. Towards a policy language for managing inconsistency in multi-context systems. In *Workshop on Logic-based Interpretation of Context: Modelling and Applications*, pages 23–35, 2011.
13. T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artif. Intell.*, 57(2-3):227–270, 1992.
14. M. Fang. Maintaining integrity constraints in semantic web. Ph.D. thesis, Georgia State University, 2013.
15. M. Fink, L. Ghionna, and A. Weinzierl. Relational information exchange and aggregation in multi-context systems. In *LPNMR*, volume 6645 of *LNCS*, pages 120–133. Springer, 2011.
16. S. Flesca, S. Greco, and E. Zumpano. Active integrity constraints. In *PPDP*, pages 98–107. ACM, 2004.
17. A. Guessoum. Abductive knowledge base updates for contextual reasoning. *J. Intell. Inf. Syst.*, 11(1):41–67, 1998.
18. Y. Guo, Z. Pan, and J. Heflin. LUBM: A benchmark for OWL knowledge base systems. *J. Web Sem.*, 3(2-3):158–182, 2005.
19. B. Motik, I. Horrocks, and U. Sattler. Bridging the gap between OWL and relational databases. *Web Semantics: Science, Services and Agents on the World Wide Web*, 7(2), 2011.
20. B. Motik and R. Rosati. Reconciling description logics and rules. *J. ACM*, 57, June 2010. Article Nr 30.
21. D. Ouyang, X. Cui, and Y. Ye. Integrity constraints in OWL ontologies based on grounded circumscription. *Frontiers of Computer Science*, 7(6):812–821, 2013.
22. P.F. Patel-Schneider and E. Franconi. Ontology constraints in incomplete and complete data. In *ISWC*, volume 7649 of *LNCS*, pages 444–459. Springer, 2012.
23. T.C. Przymusiński and H. Turner. Update by means of inference rules. *J. Log. Program.*, 30(2):125–143, 1997.
24. J. Tao, E. Sirin, J. Bao, and D.L. McGuinness. Integrity constraints in OWL. In *AAAI*. AAAI Press, 2010.
25. E. Teniente and A. Olivé. Updating knowledge bases while maintaining their consistency. *VLDB J.*, 4(2):193–241, 1995.
26. B. Thalheim. *Dependencies in Relational Databases*. Teubner-Texte zur Mathematik. B.G. Teubner, 1991.
27. M. Winslett. *Updating Logical Databases*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1990.